

Optimize Performance and Reduce Complexity

With Cisco UCS C-Series Multinode Rack Servers powered by AMD EPYC Processors



High performance



High core density



Reduce complexity



Manage locally or from the cloud

High-intensity clustered workloads need more

More organizations than ever need to deploy clustered workloads where high core density is essential. You need more cores in a rack. And more memory and I/O performance. Our new Cisco UCS® C-Series Multinode Rack Servers, powered by AMD EPYC™ processors, take these workloads out of silos and into your data center.

The result? Greater benefit from your infrastructure because you can quickly and easily reconfigure and repurpose clusters for different users at different times. You get more out of your data because we deliver 128 percent more core density than our own most dense rack servers and 71 percent more than our blade servers.

Target workloads

Supported by a new multinode rack server chassis and server nodes, this solution is designed for: service providers building cloud platforms; manufacturers simulating new designs; retailers analyzing consumer trends; compute-intensive web and gaming back-end processing; and for data scientists analyzing financial markets. Virtually any activity where data and infrastructure is critical to gaining a competitive edge.

Cisco UCS
C-Series Multinode
Rack Servers
with AMD EPYC
processors



Multinode density

- **Cisco UCS® C4200 Series Rack Server Chassis** hosts four rack server nodes in 2 rack units (2RU) with shared power and cooling.
- **Cisco UCS C125 M5 Rack Server Node** is powered by AMD EPYC™ processors that deliver the most core density in the industry.
- **Increase density** with more cores per rack.
- **Reduce complexity** with unified management.

Multinode rack servers

Our density-optimized solution is based on the Cisco UCS C4200 Series Rack Server Chassis and the Cisco UCS C125 M5 Rack Server Node. The rack server chassis provides:



Cisco UCS C4200 Series Rack Server Chassis Cisco UCS C125 M5 Rack Server Node

- **Support for four nodes.** Up to four Cisco UCS rack server nodes can populate this 2RU chassis.
- **256 cores per chassis.** With the AMD EPYC processor, industry-leading core density is achieved.
- **Shared power and cooling.** N+1 redundant power supplies support all four servers and provide sufficient power to support top-bin processors.
- **24 small-form-factor (SFF) drives.** The drive bays are allocated so that each rack server node has access to six SAS, SATA, or up to four disk drives and two NVMe drives.

Changing landscape

Workloads that use high-density clustered infrastructure were once the domain of only a select few applications. These clusters were purpose built and hosted in specialized data centers, creating islands of technology, policy, and management that are unacceptable to most IT organizations today.

The value of data

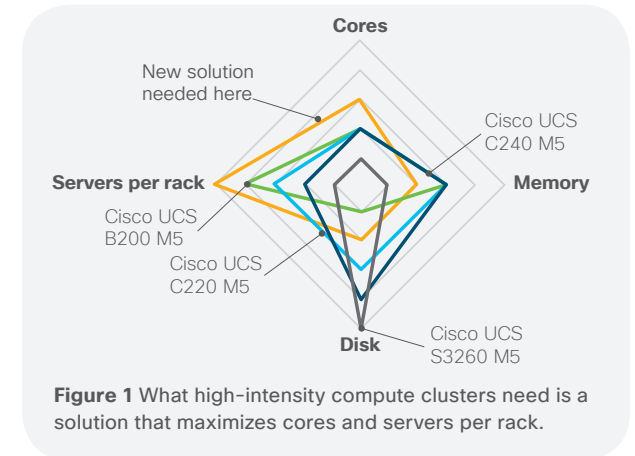
As the value of data is being recognized in nearly every industry, there is a rising demand for transforming raw data into informed business decisions. Most of today's transformational software applications are designed for scale-out computing with built-in data redundancy and fail-in-place resiliency.

Other applications use similar clusters. The major cloud providers use custom hardware to maximize density and Tier 2 providers need to emulate this model with commercial hardware in order to increase cost effectiveness. Back-end processing for web applications such as mobile gaming use massive parallelism and must also scale quickly, easily, and reliably.

The value of density

Clustered workloads need more density with more disk storage, and with more flexible I/O options compared to traditional servers.

- **Blade servers** offer density but have limited storage capacity and less I/O flexibility.



- **Rack servers** offer more disk storage and more I/O options but offer less density than blade servers.
- **Storage servers** offer high disk density but fewer CPU resources per rack unit.

What the market needs is a server that fits the profile of the orange polygon in Figure 1.

The value of infrastructure

Traditional high-intensity clustered infrastructure was used for a few specific workloads, and, like the mainframes of years past, it sits idle between tasks. With more demand for clustered infrastructure comes a demand for sharing through rapid configuration and reconfiguration of clusters – even using the servers for general-purpose workloads. With IT organizations recognizing the better return from shared resources, the only question is this: which investment is the right one?

Cisco UCS C125 M5 Rack Server Node

Our first rack server node delivers the highest core-per-rack density available from a commercially available multinode solution.



Cisco UCS C125 M5
Rack Server Node

- **Up to 2 AMD EPYC processors.** Choose the number of cores you need for your workload with consistent processor features, from 8 to 32 cores per CPU.
- **16 DIMM slots.** Up to 2 TB of memory.
- **2 PCIe 3.0 slots.** Use to connect to the I/O infrastructure that best suits your needs.
- **Optional fourth-generation Cisco VICs.** Complete programmability of the number and type of I/O devices, making hosting different workloads simple and easy.
- **Flexible OCP 2.0 slot.** For third-party I/O devices.
- **Internal drive connectivity.** Up to 6 SAS or SATA drives or up to 2 NVMe drives plus 4 SAS or SATA disk drives.
- **Internal SD or M.2 boot options.**

More performance, less complexity

We deliver the first multinode, density-optimized, shared infrastructure solution that scales with enterprise management.

In two rack units, we support four 2-socket servers powered by AMD EPYC processors, each with up to 32 cores for 256 cores per chassis—the highest core density in the industry. With large memory capacity, superior memory bandwidth, massive amounts of I/O capacity, and dedicated disk storage, the Cisco UCS C4200 Series Rack Server Chassis combined with Cisco UCS C125 M5 Rack Server Nodes brings more performance with less complexity to high-intensity compute clusters.

More performance

This Cisco® solution delivers more performance for scale-out computing clusters.

Greatest core density

Compared to our portfolio of blade, rack, and storage servers, Cisco UCS C-Series Multinode Rack Servers deliver the highest core density of all our offerings. 256 cores per chassis with help from the highest-density x86-architecture CPU in the industry. The choice of AMD EPYC processors is for more than just density: the processor also claims the [highest SPECrate® fp_base performance score](#) in the industry, helping to speed compute-intensive workloads.

High memory capacity and more bandwidth

The high core density needs matching memory capacity and more memory bandwidth to make compute-intensive applications run faster—and both are provided by AMD EPYC processors. More software and data in main memory reduces the impact of disk latency. More memory bandwidth moves data in and out of the CPU more rapidly.

More I/O

With AMD EPYC processors offering 128 lanes of I/O capacity, the Cisco UCS C125 M5 and future products have headroom to handle the I/O requirements of high-intensity compute clusters. The Cisco UCS C125 M5 supports two PCIe cards for general-purpose I/O. When equipped with fourth-generation Cisco UCS virtual interface cards (VICs), the very number and type of I/O devices are programmable on demand through any of the management approaches discussed below. This gives a new level of agility for configuring and reconfiguring compute clusters for specific workloads.

More disk space


The Cisco UCS C4200 Series Rack Server Chassis supports more disk storage than our most dense rack servers. To enhance performance, drives can connect directly to controllers built into the server CPU to maximize parallelism or they can be connected through an optional RAID controller to offload data integrity operations.

We deliver more

We deliver more of what you need to power scale-out clustered workloads:


71% more cores

per rack than our most dense blade servers



50% more servers

per rack than our most dense blade servers



20% more storage

per rack than our most dense rack servers



- Cisco UCS C125 M5
- Cisco UCS B200 M5
- Cisco UCS C220 M5

Less complexity

We have a long history of extending our product portfolio into new markets and simplifying operations with our unified role- and policy-based management. Since we introduced Cisco Unified Computing System™ (Cisco UCS) in 2009, no other vendor has yet established a unified management infrastructure that supports both blade and rack servers.

- **Integrate into your data center.** We bring the same operational approach to high-intensity compute clusters as our blade, rack, and storage servers. Now you can connect and manage your clusters in the same way you manage the rest of your IT infrastructure.
- **Programmable infrastructure.** With more than 100 identity, configuration, and connectivity parameters extracted from each server, you can use our management software to consistently configure hundreds of servers as easily as you manage a single one.
- **Maximize return on investment.** Programmable infrastructure means that you can easily configure, scale, and also reconfigure entire clusters through software, with no human touch. Maximize your return on investment by making them available to multiple stakeholders and put your infrastructure to work.
- **Secure multitenancy.** When you integrate your cluster through Cisco Application Centric Networking (Cisco ACI™), you can securely partition multiple workloads with network profiles that isolate tenants and applications.

The management advantage

What sets our multinode server solution ahead of the competition is what means most to you: reducing operational costs. We help reduce these costs by making administrators more effective with management options that make your clusters simple to configure, scale, and repurpose at a moment's notice.

With our management advantage, your infrastructure is programmable on demand, from how the disk drives are configured to the number and type of I/O devices that are accessible to the operating system or hypervisor. Even firmware is programmable on demand: if different workloads require different software revisions, updating an entire cluster of servers is a simple, zero-touch operation.

There are several ways to achieve unified management of all of your multinode servers:

- **Cisco Intersight™ management as a service.** This [cloud-based management approach](#) lets you configure and manage all of your blade, rack, storage, and multinode servers through a single interface, regardless of where your servers are installed. This role- and policy-based management service enables you to create Cisco UCS service profiles and templates that have scope across your organization worldwide. Your high-level administrators define your organization's intent for how many of your servers should be configured;

For more information

Visit cisco.com/go/ucs4200

the rest of your staff can then use these profiles to create server instances that match your intent, align with your best practices, comply with regulatory standards, and have zero configuration creep. Cisco Intersight management includes a recommendation engine that helps you proactively address problems and easily integrate with the Cisco Technical Assistance Center (TAC). This management software includes all of the capabilities of Cisco UCS Management and the Cisco Integrated Management Controller (IMC) Supervisor (see description below).

- **Cisco UCS management.** This includes a [suite of tools](#) to manage your infrastructure. Cisco UCS Manager is embedded in Cisco UCS fabric interconnects and allows your infrastructure to be configured using Cisco UCS service profiles. The software is accessible through an intuitive HTML 5 GUI, command-line interface (CLI), and the Cisco UCS Management API. The management API allows your own software to script infrastructure deployment through Microsoft PowerShell, Python, and Ruby languages. The API is also used through Puppet, Chef, Ansible, Docker, Kubernetes, and Cisco Container Platform software.
- **Cisco Integrated Management Controller (IMC) Supervisor.** The [Cisco IMC](#) provides low-level management capabilities through each server's management console. Used when Cisco UCS C-Series Servers are deployed as standalone systems, the Cisco IMC provides an interface for managing server identity, configuration, and Cisco UCS

VIC connectivity. The [Cisco IMC Supervisor](#) aggregates the consoles for multiple servers and provides a simplified interface for configuring large numbers of servers such as those in a high-intensity compute cluster.

Gain more value from your data and infrastructure

We help you increase the value you gain from your data by providing a high-performance, high-density solution for a growing category of applications. We help you increase the value of your infrastructure by bringing high-intensity clustered workloads out of silos and into your data center.

This is an investment in the future, because as the use cases for this approach expand, so does your capability to treat compute-intensive workloads just like other workloads in your data center. With unified management across our entire product portfolio, you maximize the value of your administrative staff and make vast amounts of computing resources available to the stakeholders you serve.